Towards a Theory of Self-organization

E. Anceaume[1], X. Defago[2], M. Gradinariu[1], and M. Roy[3]

[1] IRISA, Rennes, France
{anceaume, mgradina}@irisa.fr
[2] JAIST and PRESTO, JST, Japan
defago@jaist.ac.jp
[3] LAAS-CNRS, Toulouse, France
roy@laas.fr

Self-organization is an evolutionary process in which the effects of the environment are minimal; i.e., where the development of new, complex structures primarily takes place in and throughout the system itself. Natural phenomena, living forms, or social systems (e.g., growing crystals, cells aggregation, ant colonies) are examples of self-organizing systems in which a global order of the system emerges from local interactions. In the newly emerging fields of distributed systems (p2p, ad-hoc networks, sensor networks, cooperative robotics), self-organization has become one of the most desired properties. The major feature of all recent scalable distributed systems is their extreme dynamism in terms of structure, content, and load. In peer-to-peer systems, self-organization is handled through protocols for node arrival and departure, based either on a fault-tolerant overlay network, such as in CAN, Chord, Pastry, or on a localization and routing infrastructure [2]. In ad-hoc networks, self-organizing solutions have been designed to cluster ad-hoc nodes [4]. Self-organizing algorithms have also been developed to arrange mobile robots into predefined geometric patterns (e.g., [3]).

Informal definitions for self-organization, or the related self* properties (e.g., self-configuration, self-healing or self-reconfiguration) have been proposed previously [4]. Zhang and Arora [4] propose the concepts of self-healing and self-configuration in wireless ad-hoc networks.

The correctness proofs for self-organizing systems should be based on a well-founded theoretical model that can capture the dynamic behavior of these systems. Hence, the characterization of the self-organizing aspects of these systems cannot solely focus on the non-dynamic periods since they may be absent or very short. Moreover, defining self-organization as a simple convergence process towards a stable predefined set of admissible configurations is inadequate for two reasons. First, it may be impossible to clearly characterize the set of admissible configurations since, in dynamic systems, a configuration should include the state of some key parameters that have a strong influence on the dynamicity of the system. These parameters can seldom be quantified *a priori* (e.g., the status of batteries in sensor networks, or the data stored within p2p systems). Second, due to the dynamic behavior of nodes, it may happen that no execution of the system converges to one of the predefined admissible configurations.

P. Fraigniaud (Ed.): DISC 2005, LNCS 3724, pp. 505–506, 2005.

We propose a formal specification of the self-organization notion. Our specification is based on the locality principle, i.e., the fact that interactions and knowledge are both limited in scope. We formalize this idea [1], leading first to the notion of *local self-organization*. Intuitively, a locally self-organizing system should reduce the entropy of the system in the neighborhood of each node. For example, a locally self-organizing p2p system forces components to be adjacent to components that improve, or at least maintain, some property or evaluation criterion. We then formalize the notion of *self-organization* by imposing the system to be locally self-organizing at all its nodes, and by ensuring that despite its dynamicity, the system entropy progressively decreases. That is, self-organization strongly relies on the local self-organization property, and on the behavior of the system during the connection/disconnection actions. According to this behavior, the system guarantees different levels of self-organization, namely, the weak and the strong self-organization. The weak self-organization is defined in terms of two properties. The *weak liveness* property says that either (1) infinitely often, there are static fragments, i.e., sequences of configurations with no connections/disconnections, during which the knowledge of the system enriches, or (2) all processes have reached a stable state. The *safety* property states that, during all static fragments, system knowledge does not decrease. The weak self-organization definition applies to static fragments. Nothing is guaranteed during dynamic ones (i.e., sequences of configurations in which connections/disconnections occur). For instance, the weak liveness does not forbid processes to reset their neighbors lists after each connection/disconnection. To prevent the system from "collapsing" during dynamic fragments, *strong self-organization* specifies a stronger property guaranteeing that for all the processes of which the neighborhood is unchanged, information is maintained. Specifically, this ensures the existence of a non-empty group of processes for which local information has been maintained between the end of a static fragment and the beginning of the subsequent one. The second contribution of this work is a case study. Using our framework we prove the weak self-organization of Pastry and CAN, two well known peer-to-peer overlays.

Future investigation focuses on the design of a probabilistic extension of our model motivated by the fact that connection/disconnection actions are well-modeled by probabilistic laws.

References

1. E. Anceaume, X. Defago, M. Gradinariu, and M. Roy. Towards a theory of self-organization. Technical Report 1694, IRISA, 2005.
2. Stephanos Androutsellis-Theotokis and Diomidis Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Comput. Surv.*, 36(4):335–371, 2004.
3. I. Suzuki and M. Yamashita. Distributed anonymous mobile robots: formation of geometric paterns. *SIAM Journal of Computing*, 28:1347–1363, 1999.
4. H. Zhang and A. Arora. Gs3 : Scalable self-configuration and self-healing in wireless networks. *Proc. of the 21st Annual ACM Symposium on Principles of Distributed Computing (PODC'02)*, pages 58–67, 2002.